Behavioral/Cognitive

The Anterior Cingulate Gyrus Signals the Net Value of Others' Rewards

Matthew A. J. Apps[1,2,3] **and Narender Ramnani**[1]

[1]Department of Psychology, Royal Holloway, University of London, Egham TW20 0EX, United Kingdom, [2]Nuffield Department of Clinical Neuroscience, University of Oxford, John Radcliffe Hospital, Oxford OC3 9DU, United Kingdom and [3]Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, United Kingdom

Evaluating the costs and benefits of our own choices is central to most forms of decision-making and its mechanisms in the brain are becoming increasingly well understood. To interact successfully in social environments, it is also essential to monitor the rewards that others receive. Previous studies in nonhuman primates have found neurons in the anterior cingulate cortex (ACC) that signal the net value (benefit minus cost) of rewards that will be received oneself and also neurons that signal when a reward will be received by someone else. However, little is understood about the way in which the human brain engages in cost–benefit analyses during social interactions. Does the ACC signal the net value (the benefits minus the costs) of rewards that others will receive? Here, using fMRI, we examined activity time locked to cues that signaled the anticipated reward magnitude (benefit) to be gained and the level of effort (cost) to be incurred either by a subject themselves or by a social confederate. We investigated whether activity in the ACC covaries with the net value of rewards that someone else will receive when that person is required to exert effort for the reward. We show that, although activation in the sulcus of the ACC signaled the costs on all trials, gyral ACC (ACC$_g$) activity varied parametrically only with the net value of rewards gained by others. These results suggest that the ACC$_g$ plays an important role in signaling cost–benefit information by signaling the value of others' rewards during social interactions.

Key words: cingulate; effort; empathy; reward; social; value

Introduction

Theories of decision making highlight that choices are made by weighing up rewarding benefits against the costs incurred to receive desired outcomes (Bautista et al., 1998; Phillips et al., 2007). Such cost–benefit analyses underpin decision making in many species, including humans (Charnov, 1976; Kagel and Levin, 1986; Bautista et al., 2001). However, it has often been overlooked that such decisions are not made in a social vacuum (Walton and Baudonnat, 2012). Our decisions are influenced by cost–benefit analyses that we apply to understand others' rewarding outcomes. Here, we investigate the neural processes that are involved in evaluating the costs and benefits of rewards received oneself and by others.

The anterior cingulate cortex (ACC) is engaged when processing information about rewards (Rogers et al., 2004; Sallet et al., 2007) and there is evidence that it plays an important role in

cost–benefit evaluation (Phillips et al., 2007; Walton and Baudonnat, 2012). Neurons in the ACC encode the net value of rewards (cost–benefit) at the time of cues that are instructive of how much effort is required (cost) for primary reinforcement (benefit) (Kennerley et al., 2009; Kennerley et al., 2011; Hillman and Bilkey, 2012). Neuroimaging studies have shown that activity in the ACC at the time of such cues is a function of the magnitude of a secondary reinforcer and the anticipated amount of effort (Botvinick et al., 2009; Croxson et al., 2009). This evidence suggests that the ACC signals the net value of rewards when they are to be received oneself.

The ACC is also implicated in the processing of social information (Behrens et al., 2009). Within the ACC, there is a dissociation between information processing in the sulcus (ACC$_s$) and the gyrus (ACC$_g$). Lesions to the ACC$_s$ that leave the ACC$_g$ intact disrupt first person decision making (Kennerley et al., 2006). In contrast, lesions to the ACC$_g$ that leave the ACC$_s$ intact disrupt social behavior and the processing of social stimuli (Rudebeck et al., 2006). However, there is evidence that the ACC$_g$ and ACC$_s$ are both sensitive to reward-related information during decision making (Behrens et al., 2009). Although the ACC$_s$ processes information about one's own rewarding outcomes, the ACC$_g$ is engaged when monitoring the outcomes of others' choices or the outcomes of choices that will be experienced by others (Behrens et al., 2008; Apps et al., 2013a; Chang et al., 2013). This implicates the ACC$_g$ as a candidate for processing the net value of rewards that others will receive (Apps et al., 2013b). However, no previous

study has investigated activity in the brain when subjects weigh up the costs and benefits of rewards that others will receive.

Using fMRI, we examined activity time-locked to cues that signaled the level of economic reward available, the cost incurred for receipt, and also whether the rewards and costs pertained to the first person or to a third person. We tested the hypothesis that ACC_g activity signals the anticipated net value of rewards to be received by a third person when they are required to exert effort to gain them and that ACC_s activity signals net value in relation to rewards gained oneself.

Materials and Methods

Subjects

Subjects were 16 healthy, right-handed participants screened for neurological disorders (age 18–32 years; 13 female). Two subjects were excluded from the analyses. Both subjects failed to maintain a belief in the deception and one of these subjects failed to perform the judgment task (see Judgement task, below) better than chance (one male). All participants gave written informed consent. The studies were approved by the Royal Holloway University of London Psychology Department Ethics Committee and conformed to the regulations set out in the CUBIC MRI rules of Operations. Subjects were paired up with one of two confederate participants, whom they believed were also naive participants. The subjects believed that they would be paid for their participation based on their performance of the task during a scanning session (see Effort task and Judgement task, below). They also believed that the confederate would be paid based on their performance in the same manner.

Apparatus

Subjects lay supine in an MRI scanner with the fingers of the right hand positioned on an MRI-compatible response box. Stimuli were projected onto a screen behind the subject and viewed in a mirror positioned above the subjects face. Presentation software (Neurobehavioral Systems) was used for experimental control (stimulus presentation and response collection). A custom-built parallel port interface connected to the presentation PC received transistor-transistor logic (TTL) pulse inputs from the response keypad, as well as TTL pulses from the MRI scanner at the onset of each volume acquisition, allowing events in the experiment to become precisely synchronized with the onset of each scan. The timings of all events in the experiment were sampled accurately, continuously, and simultaneously (independently of presentation) at a frequency of 1 kHz using an A/D 1401 unit (Cambridge Electronic Design). Spike2 software was used to create a temporal record of these events. Event timings were prepared for subsequent general linear model (GLM) analysis of fMRI data (see Statistical analysis, below).

Experimental design

The aim of this experiment was to examine the processing of cues that instructed a first person and a third person as to how much reward they would receive after the exertion of differing levels of effort. Subjects performed a task over 2 d with a training partner (confederate). On the first day, the subject and the confederate learned the associations between a set of instruction cues, a financial reward, and how much effort they were required to expend for its receipt. On the second day, both agents continued to perform effortful actions to receive rewards. During this session, the subject performed these trials while inside the MRI scanner with the training partner situated in the adjacent control room.

A 2 × 2 × 2 factorial design was used to examine activity time locked to instruction cues (Fig. 1). The first factor was agency. On each trial, either the subject (first person) or the confederate (third person) performed a series of cued button presses (or "cancellations") on a keypad to receive a reward. The second factor was the reward level that was obtainable on each trial. This could be either high reward (HR) if 16 UK pence (p) was obtainable on the trial or low reward (LR) if only 4 p was obtainable. The third factor was the level of effort. There were four levels of effort (two, three, eight, or 12 responses), which corresponded to the number of cancellations (cued button presses) that were required to

receive the reward. These were collapsed into either low effort (LE) for two or three cancellations or high effort (HE) for 8 or 12 cancellations for the factorial design. All cues were color coded on each trial such that the first person responded when stimuli were blue and the third person when stimuli were brown. On each trial, the instruction cues signaled the level of reward available and the effort required by either the first person or the third person. The instruction cue stimuli were based on those used in previous studies (Knutson and Bossaerts, 2007; Croxson et al., 2009) investigating first-person reward prediction processing. The stimuli were 80-mm-diameter circles containing crosshairs. The position of the crosshairs indicated both the amount of reward that was obtainable and the number of cancellations required to receive that reward. Reward was represented vertically on the circle (16 p was high on the circle, 4 p was low). Effort was represented horizontally with increasing levels of effort represented from left to right.

In total, there were 16 different trial types dependent on the reward level, effort level, and the agent performing the cancellations. There were eight different trial types for each level of agent (Fig. 1).

Trial structure

Each trial (Fig. 1) began with one of 16 different color-coded instruction cues. These cues indicated both the level of reward that was available on each trial and also the level of effort required for its receipt. The color of these cues also indicated who would have to perform the cancellations on each trial (blue for the first person, brown for the third person). After the instruction cue, there was an effort period during which cued button presses were performed on a keypad. During the effort period on the first-person trials, subjects were required to make a series of cued button presses (cancellations). On the third-person trials, the cancellations were actually preprogrammed computer controlled responses (see Judgement task, below). At the end of the effort period, a stimulus then displayed the number of cancellations that had been made during the effort period. After this stimulus, a trigger cue (3 lines with 16 p over the left line, 4 p over the middle line, and 0 p over the right line) was presented on the screen, which cued the first person or the third person to make a judgment of the amount of reward that would be received by the other agent on that trial. Each line corresponded to one button on the keypad. Subjects had 750 ms to make their response. If they did not respond in this time window, it was classified as an incorrect response. After this, a feedback cue indicated the accuracy of the judgment ("correct" if the judgment was correct and "−10 p" if incorrect) and then a feedback cue indicated the reward received by the agent who performed the cancellations (16, 4, or 0 p). In total, there were 192 trials, 96 first-person trials in which the subjects made cancellations and 96 in which they monitored the third person's cancellations. Each instruction cue was presented on 12 trials.

Task

Subjects performed two tasks during scanning. On first-person trials, subjects performed an "effort task," in which cancellations were made to receive financial rewards. On the third-person trials, subjects performed a judgment task, monitoring the third person's performance of the effort task and indicating the amount of reward they would receive.

Effort task

During scanning, subjects performed trials in which they were required to make the correct number of cancellations to receive a financial reward. The effort task required subjects to make a series of cancellations during the effort period. Each cancellation was cued by the position of a square stimulus above one of four lines on the screen. Each line on this cue corresponded to one button on the keypad. The position of the square highlighted a target button. A cancellation constituted one press of the target button: one finger movement of one finger on the right hand. Once this target button was pressed, the position of the square would move to highlight a new target button. Each target button was always different from the previous. Subjects could make up to 14 of these cancellations during the fixed time window of the effort period (6600 ms). Subjects were only rewarded on a trial if they performed exactly the number of cancellations specified by the cue. If they performed more
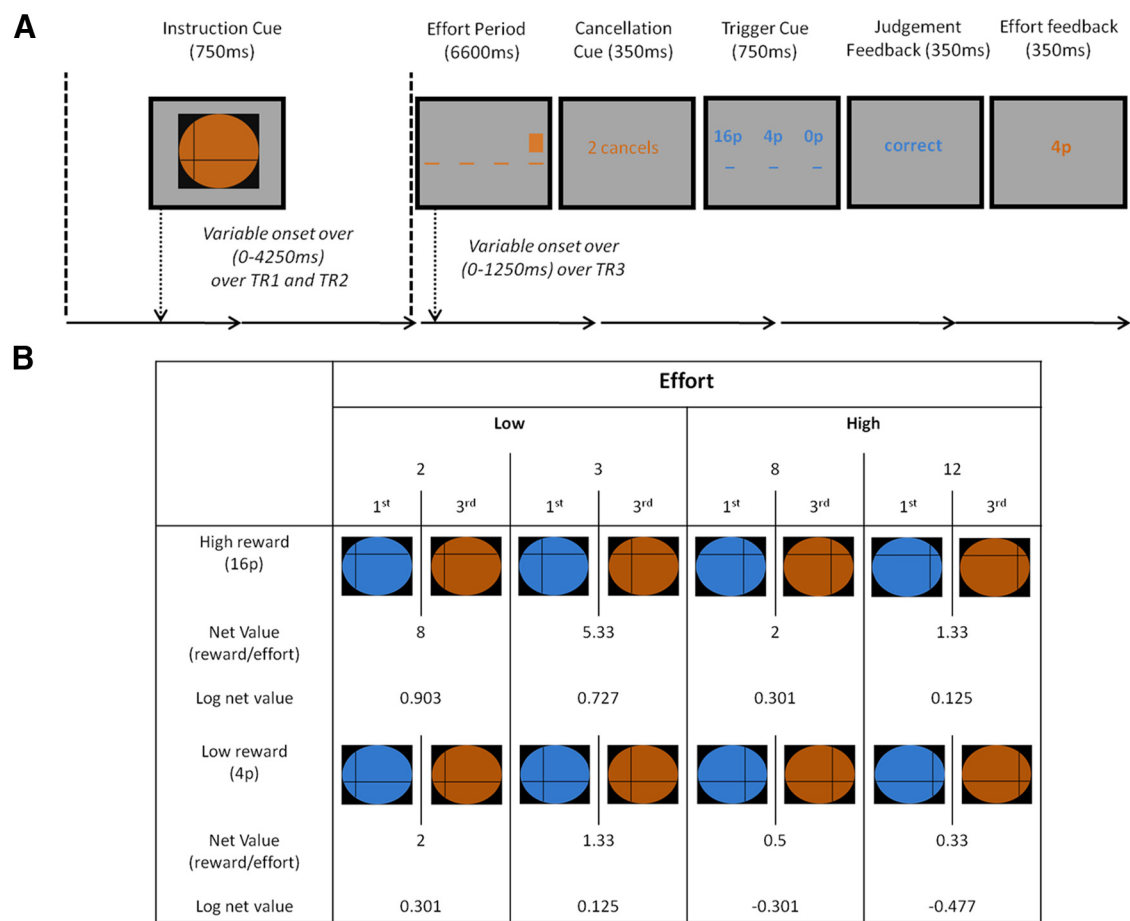
Figure 1. *A*, Trial structure. Participants performed trials that began with a color-coded instruction cue. Blue indicated that the subjects (first person) would perform button presses to receive a financial reward and brown indicated that the button presses would be performed by a confederate (third person). The position of crosshairs on these stimuli indicated the level of reward available (16 or 4 p) and the number of button presses required for reward receipt (2, 3, 8, or 12). After this was the effort period, in which the required button presses were made by cancelling out highlighted squares that each corresponded to one button on a keypad. After this was a cue that indicated the number of button presses ("cancellation cue") and then a trigger cue. At the time of the trigger cue, on the third-person trials, the subjects were required to indicate how much reward the third person would receive (16, 4, or 0 p if they made the incorrect number of button presses). After this, there was feedback for performance on this judgment task and, finally, there was feedback for the individual who performed the effort task. The instruction cue onsets were jittered over the first two scans (TRs) of each trial to sample evoked hemodynamic responses time locked to these events evenly and independently from the other elements in the trials. The dotted lines indicate the cue that was jittered over the first two TRs. *B*, Experimental design displayed in a table showing the 16 different conditions in the experiment. Net value was calculated as the level of reward divided by the level of effort. These were then log transformed to create the parameters used for the parametric analysis.

cancellations or fewer cancellations, they would not obtain the rewarding outcome. Using such an approach ensured that the level of effort expended by participants on the task was closely matched to that specified by the cue. The order of cancellations was randomized across the experiment and within each effort period. Subjects were therefore unable to make any prediction about which button would be the next target. Subjects were instructed to make cancellation responses as quickly and accurately as possible for every level of effort.

We used this cancellation task as a corollary of effortful exertion. A large number of previous studies have shown that the number of actions that are performed is in index of effort. Such studies have shown that the number of lever presses or the number of actions or button presses modulates reaction times and choice behavior in rats, humans, and monkeys. Reaction times to reward predicting stimuli increase as the number of actions increases. Choices are made in favor of rewards associated with fewer actions than the same magnitude of reward associated with more actions (Walton et al., 2006; Phillips et al., 2007; Salamone et al., 2007; Floresco et al., 2008a; Floresco et al., 2008b; Croxson et al., 2009; Floresco and Whelan, 2009; Kennerley and Wallis, 2009; Kennerley et al., 2009; Walton et al., 2009; Day et al., 2010; Nunes et al., 2010). This supports the argument that the number of actions that are to be performed modulates the value of the rewarding outcome as a result of the costly nature of increased effortful exertion. Therefore, by using a task in which the num-

ber of actions was the index of effortful exertion, we have used a paradigm that is well known to modulate reward desirability.

In our task, we assumed that effort could be equated to the number of cancellations made, so it was therefore important that the number of required cancellations as specified in the instruction cue was closely related to the actual amount of actions performed by the subjects. To do so, we first ensured that the task was very simple and the subject and confederate were overtrained together on the task before they entered the scanner during a training phase. As a result, subjects would not press the incorrect target button for a cancellation on many trials. Therefore, the number of cancellations made would be very closely related to the number of button presses made and therefore the effort expended. Furthermore the overtrained confederate was instructed to make very few mistakes in terms of which button needed to be pressed. As a result, the subject learned to expect that the confederate would make exactly the same number of button presses as required to make the correct number of cancellations.

Second, in the cancellation task, the subject was required to perform the correct button press for a new cancellation cue to be presented. If they pressed the incorrect button, then the cancellation cue would not move to a new location. Therefore, the subject inside the scanner was able to monitor the number of new targets presented to infer the number of correct button presses being made by the confederate. These two features

of the design ensured that the subject was able to know the number of cancellations on every trial and also infer that this would be highly correlated with the actual number of button presses and therefore the level of effortful exertion.

During scanning, subjects were told that they were accumulating monetary rewards for their performance on this task. Therefore, subjects believed that they were earning the reward available on each first-person trial if they performed the correct number of cancellations. Subjects were told that if they performed every cancellation correctly, they would accumulate £10 as payment for the experiment. However, unbeknown to the subjects, they would be paid £10 for participation regardless of their task performance.

Judgment task
In addition to the effort task, subjects also performed a judgment task on the trials in which the third person was performing the effort task. For this task, the subjects were required to indicate the level of reward that would be received by the third person, which could be 16 or 4 p for the correct number of cancellations or 0 p if the number of cancellations was incorrect. Subjects were required to perform this judgment on every trial performed by the third person. Subjects believed that they were punished for each incorrect judgment (when the subject indicated that the amount of money earned by the confederate on the third-person trials was different from the amount they would actually earn) by 10 p being removed from the money that they were accruing on the effort task. A correct judgment left the rewards accumulated during the effort task the same. Therefore, subjects believed that if they performed every set of cancellations correctly, but every judgment incorrectly, they would receive no payment for the experiment. Therefore, subjects were motivated to perform both tasks to the same degree of accuracy. This punishment ensured that subjects attended to the rewarding value and the effort information contained in the instruction cues on the third-person trials. Importantly, the punishment used as the motivation for the subject on the third-person trials was unrelated to the anticipated reward and effort level that would be processed at the time of the instruction cues.

This task required subjects to monitor both the amount of effort and the reward magnitude in the instruction cue, but also the number of cancellations performed. To perform this task correctly, subjects needed to monitor the cancellations made during the effort period. They could do so in two ways. First, subjects could monitor the number of times that cues disappeared during the effort period. Second, subjects were explicitly instructed as to how many cancellations had been made by the confederate at the end of the effort period by an additional cancellation cue. As a result, subjects were able to monitor the number of cancellations made by the confederate without difficulty.

To maintain experimental control, we deceived participants as to the nature of the third person. Although subjects believed that they were performing the task with another real participant the responses they saw were computer generated. This approach was necessary to maintain control over the performance of the third person. However, to ensure that payment for the experiment was ethical and not dependent on the third person's performance, all subjects were paid the same amount for participation. Subjects were thoroughly debriefed using a standard set of questions described previously (Ramnani and Miall, 2004; Apps et al., 2012; Apps et al., 2013a). Only two participants, who were excluded from the analysis, showed any awareness of the nature of the deception.

The effort task used in this study is very similar to that used by Croxson et al. (2009) to investigate first-person effort discounting. Given this similarity, it is important to note the differences between the task used here and that used by Croxson et al. (2009). There are two important differences between the effort task used in that study and that used here. First, unlike in the present study, there were no constraints placed upon the time that subjects had to make cancellations. Second, also unlike the present study, subjects were only presented with the correct number of targets to be cancelled. However, these two aspects of their task were not suitable for the purposes of this study. A crucial aspect of our design was that subjects were required to make a judgment on the reward to be received by a third person. This task ensured that subjects attended to the effort and reward levels at the time of the instruction cues on the third-

person trials. Without a temporal constraint on the effort period, there would be no possibility of making an incorrect number of cancellations, so the confederate would not make errors on the effort task. Without confederate errors, the subject could perform the judgment task by attending to the level of reward at the time of the instruction cues on the third-person trials and not the level of effort. Therefore, in this experiment, a temporal window was a necessity. In addition, in this experiment, subjects could cancel up to 14 targets, more than the maximum instructed number of 12, regardless of how many cancellations they were required to make. This created the potential for catch trials in which the confederate made an error in the number of cancellations, which the subject would need to identify correctly to maximize their own financial rewards and to perform the judgment task correctly. These two distinctions from the task used by Croxson et al. (2009) ensured that subjects attended to both the effort and reward level on every confederate trial.

Procedure
Training
Subjects were trained in 2 phases 1 d before scanning. In the first phase, the subject was seated in front of a monitor with a confederate (third person). They were each provided with a response keypad. Both the confederate and the subject performed the effort task on separate trials. During this session, both the confederate and subject learned the contingency between the position of the crosshairs on the instruction cue stimulus, the amount of reward (16 or 4 p), and the required number of cancellations (button presses) to receive the reward. They were informed before this that there would be two levels of reward and that they would have to make two, three, eight, or 12 button presses. During training, there were 64 first-person trials, in which the subject performed the cancellations and 64 third-person trials performed by the confederate. The subjects were told that the rewards were fictional during training and that their payment for the experiment would be based solely on performance during the scanning session.

In this session, as the subjects were seated next to the confederate, the confederates performed the effort task on separate trials from the subject. Because the confederates were paired with multiple different subjects throughout the piloting and experimental phases, they were highly overtrained on the effort task. To ensure that subjects maintained the belief that the confederates were naive participants like themselves, they were told to make deliberate errors in the number of cancellations performed during the first phase of training to mimic the learning of a real participant.

In the second phase of training, subjects practiced the task that would be performed during the scanning session (see Scanning session below). The subject performed this from inside a mock scanner, with the confederate seated in front of a monitor adjacent to the mock scanner. The subject was played the sound of a genuine scanner's EPI sequence via headphones. During this training phase and during scanning, the responses on the third-person trials were computer controlled.

Scanning session
Before scanning, subjects were shown the confederate seated in front of a monitor in the control room next to the scanner. They were told that they would see all of the responses of the third person in real-time inside the scanner. In fact, these responses were all computer-controlled, preprogrammed responses. The apparent reaction times of the confederate during the effort task were pseudorandomly organized. The reaction times of the second to 12th button presses fitted a normal distribution around a mean (525 ms), with a range of 325–725 ms. The confederate's reaction times to the first target were extended to reflect the unpredictability of the onset of this target. These formed a normal distribution around a mean of 600 ms, with a range of 400–800 ms. These timings were based on the reaction times of five participants during a pilot experiment. The apparent reaction times of the confederate were programmed to not be different regardless of the number of cancellations to be made on the trial so the confederate behavior appeared to conform to the instruction responding as quickly and accurately as possible. It was also noted that such behavior was exhibited by subjects in a pilot experiment in which no difference was found between reaction times of the first two button

presses in the two-cancellation conditions compared with the 12-cancellation conditions.

Key for the design of this experiment was that subjects attended to both their own instruction cues and those of the third person in the same manner. There was one potential caveat to the judgment task used to motivate subjects to attend to the instruction cues of confederate. Specifically, if the confederate performed the correct number of cancellations on every trial, the subject could, over time, learn to perform the judgment without attending to the level of effort, only the reward level. To address this potential confound, errors were preprogrammed into the behavior of the confederate. On nine of the trials in which the effort task was performed by the confederate, the number of cancellations performed was not correct for the instruction cue presented. These "catch" trials were used as an index of the extent to which subjects were attending to the effort expended by the confederate.

Behavioral analysis

Behavioral analyses were performed in SPSS version 16. Performance on the effort task was analyzed by performing a repeated-measures ANOVA on the effect of effort on task accuracy. Planned pairwise comparisons were then performed between the two- and 12-button-press conditions to ensure that there was no significant effect of the number of button presses (i.e., difficulty) on performance of the task. For the judgment task, paired-samples t tests were performed to examine the difference between the accuracy on the task (i.e., judging the correct reward level on the trial) and chancel level (33%). Trials in which the subject failed to respond within the 750 ms response window were included as errors. In addition, we performed two ANOVAS on the accuracy on the judgment task, one that looked for an effect of reward level or effort level on accuracy and a second that looked for an effect of net value on accuracy.

Functional imaging and analysis
Data acquisition

T1-weighted structural images were acquired at a resolution of $1 \times 1 \times 1$ mm using an MPRAGE sequence. A total of 1164 EPI scans were acquired from each participant. Thirty-four slices were acquired in an ascending manner at an oblique angle ($\approx 30°$) to the AC-PC line to decrease the impact of susceptibility artifact in subgenual cortex (Deichmann et al., 2003). A voxel size of $3 \times 3 \times 3$ mm (25% slice gap, 0.8 mm) was used; TR = 2.5 s, TE = 32, flip angle = 81°. The functional sequence lasted 48.5 min. Immediately after the functional sequence, phase and magnitude maps were collected using a GRE field map sequence ($TE_1 = 5.19$ ms; $TE_2 = 7.65$ ms).

Image preprocessing

Scans were preprocessed using SPM8 (www.fil.ion.ucl.ac.uk/spm). The EPI images from each subject were corrected for distortions caused by susceptibility-induced field inhomogeneities using the FieldMap toolbox (Andersson et al., 2001). This approach corrects for both static distortions and changes in these distortions attributable to head motion (Hutton et al., 2002). The static distortions were calculated using the phase and magnitude field maps acquired after the EPI sequence. The EPI images were then realigned and coregistered to the subject's own anatomical image. The structural image was processed using a unified segmentation procedure combining segmentation, bias correction, and spatial normalization to the MNI template (Ashburner and Friston, 2005); the same normalization parameters were then used to normalize the EPI images. Last, a Gaussian kernel of 8 mm FWHM was applied to spatially smooth the images to conform to the assumptions of the GLM implemented in SPM8.

Statistical analysis
First-level analyses

First-level GLMs were created for both factorial and parametric analyses.

Factorial analysis. There were 10 event types. Each event type was used to construct a regressor by convolving the stimulus timings with the canonical HRF. Each of the eight conditions was modeled as a separate regressor. In addition, one regressor modeled the activity during the

effort periods (regardless of whether it was a first-person or third-person trial) and another regressor modeled the onsets of the other trial elements on every trial. Trials in which the subject failed to perform the correct number of cancellations during the effort period, failed to respond within 750 ms of the onset of the trigger cue for the judgment task, or failed to make the correct response on the judgment task were modeled separately as an extra regressor. This regressor included the onsets from all of the trial elements from missed trials. The residual effects of head motion were modeled as covariates of no interest in the analysis by including the six head motion parameters estimated during realignment.

Parametric analysis. Two GLMs were created at the first level and used a parametric approach. Each of these GLMs was constructed using the same events as those used in the factorial analysis. For these GLMs, however, the instruction cue regressors were collapsed down into one regressor for the first-person instruction cues and one regressor for the third-person instruction cues. To create parametric regressors, we divided the reward magnitude by the number of cancellations required and log transformed these values, as described previously (Croxson et al. (2009). The parameters outlined in Figure 1 (the log-transformed net values) were used as first-order parametric modulators of first-person and third-person instruction cue events. In addition, we included additional parametric modulators that were scaled with the effort level. To examine activity that varied with net value, the net value parameters were orthogonalized with respect to the effort parameters. This ensured that activity that varied with first-person net reward values, third-person net reward values, or both could not be explained by the level of effort alone. The second GLM was similar except that the effort parameters were orthogonalized with respect to the net value parameters.

Second-level analysis

Random-effects analyses (full-factorial ANOVAs) were applied to determine voxels significantly different at the group level. SPM{t} contrast images from all subjects at the first-level were input into second-level full factorial design matrices. F-contrasts were conducted in each of the second-level random-effects analyses. For the whole-brain analyses, FDR correction was applied. To test the specific hypotheses in the ACC, 80% probability masks of the ACC_g and ACC_s were created and used as the search volumes for small volume correction (see "Anatomical Localization"). In addition, we also used the coordinates of Croxson et al. (2009) for small volume correction because their study examined activity at the time of instruction cues that indicated to a subject the effort level required and the level of reward available to them. Small volume corrections were applied as a sphere with 8 mm radius around the peak coordinates from their analysis that looked for an interaction between reward and effort. This correction was applied by making a mask combining each of the spheres around their peak coordinates for the comparable contrast.

Anatomical localization

To correct for multiple comparisons for our main hypotheses, we used 80% probability anatomical masks of the ACC_g and ACC_s. To create each mask, subject-specific masks of the ACC_g and ACC_s were constructed in FSL (http://www.fmrib.ox.ac.uk/fsl/). Although the cytoarchitectonic boundaries of the ACC have no corresponding gross anatomical landmarks, we defined the anatomical boundaries based on the location of these boundaries in previous literature investigating cingulate cytoarchitecture (Vogt et al., 1995). We used a posterior horizontal extent to each mask that lay 22 mm posterior to the anterior commissure (i.e., the posterior border of the midcingulate cortex; Vogt et al., 1995). We included all voxels that lay within the ACC_s or the ACC_g extending anterior to this border, including subgenual cingulate cortex. The final ACC_s and ACC_g masks included only voxels that were within each region in 80% of our subjects. We defined our anatomical mask of the ACC_s as cytoarchtiectonic zones 24c, 24c′, 32, and 32′ as defined by Vogt et al. (1995), who note that when there is only a single cingulate sulcus and no paracingulate sulcus, areas 24c and 24c′ lie on the ventral bank of the sulcus and areas 32 and 32′ lie in the dorsal bank of the sulcus. When a paracingulate sulcus is present, areas 24c and 24c′ lie in the both the ventral and dorsal banks of the primary cingulate sulcus and
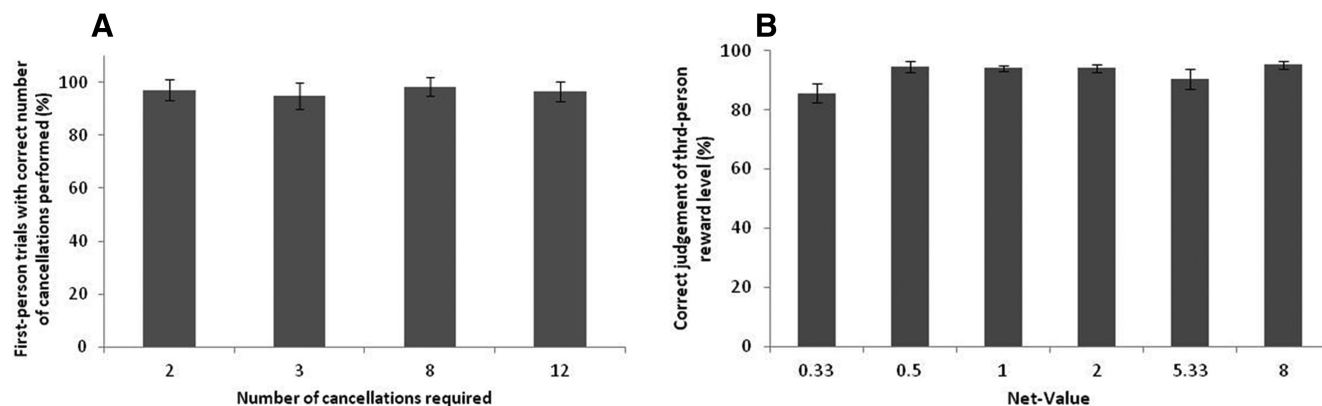
A   **B**

Figure 2. Behavioral results. **A**, Behavioral results for the performance of the effort task by the subjects themselves. Behavioral performance was high for each effort level and showed no significant differences between effort conditions. **B**, Performance at judging the level of reward that was being received by the confederate for each of the six net value levels. A repeated-measures ANOVA showed no difference in performance by net value. There was also no effect of effort or reward on the task accuracy.

areas 32 and 32′ lie on the additional paracingulate gyrus and extend over the ventral bank of the paracingulate sulcus. We created masks of the ACC$_s$ using exactly the same anatomical criteria. When there was a single cingulate sulcus, the mask covered the dorsal and ventral banks of the sulcus. When there was an additional paracingulate sulcus, the mask included both the dorsal and ventral banks of the cingulate sulcus and extended up to and including the ventral bank of the paracingulate sulcus.

Results

Behavioral results

The subjects performed two tasks while inside the MRI scanner. On first-person trials, they performed button presses to receive rewards themselves. On the third-person trials (Fig. 1), they performed a second task, judging the level of reward (16, 4, or 0 p) that the confederate would receive on that trial after they had monitored the confederate's responses. For the effort task, subjects were required to make 2, 3, 8, or 12 button presses (cancelling out 1 of 4 visually cued targets by pressing 1 of 4 corresponding buttons on a keypad) to receive a financial reward (16 or 4p). These button presses were made during a 6600 ms effort period. An important issue in this experiment was that the effort task constituted effort and not difficulty. In previous studies, the effort period (Botvinick et al., 2009; Croxson et al., 2009) was not constrained by a time period, cancelling out a large number of targets was not more difficult than a small number. In this study, the fixed response window may have caused subjects to find it more difficult to complete the 12 button presses than to make 8, 3, or 2 responses. This would confound any interpretation because effort-related activity would have been confounded with activity occurring due to the probability of success on the task, which could have led to subjects making a negative reward prediction or risk-related prediction at the time of the cues on the third-person trials. However, if subjects' performance was high and consistent across the effort and reward levels, then any potential confounds would be orthogonal to first-person or third-person net value and therefore could not account for activity covarying with net value at the time of the instruction cues.

To determine whether effort was confounded with first-person effort task difficulty, the behavioral accuracy of subjects across each of the four effort levels were examined on the first-person trials (Fig. 2A). Correct response trials were those in which the subject made exactly the same number of cancellations as those specified by the cue. Notably, task accuracy was high (mean = 96.71%), suggesting that subjects did not find the task

difficult. In addition, as can be seen in Figure 2A, the accuracy was high for all four effort levels on the task. A repeated-measures ANOVA was performed examining the effect of effort on task accuracy (percentage of correct responses). No main effect of effort on task accuracy was identified ($F_{(2.18, 28.37)} = 2.098$, $p = 0.198$). A planned pairwise comparison between the 2- and 12-button press trials showed no significant difference in accuracy ($t_{(13)} = 1.528$, $p = 0.151$), indicating that they were not significantly less accurate at performing 12 button presses compared with 2 button presses. This suggests that the increased amount of effort in the task did not cause a significantly increased level of difficulty.

The second task performed by subjects was a judgment of the reward level that would be received by the third person. Subjects were required to monitor the responses of a third person (confederate) and indicate whether they would receive HR (16 p), LR (4 p), or no reward (0 p) on each trial. Performance on this task was an important index of subjects' understanding of the level of reward available and the effort necessary for its receipt on the third-person trials. It was of particular importance that subjects performed catch trials in which the third person made the incorrect number of button presses above chance level (33.3%). On these trials, subjects could not perform the judgment task correctly without attending to the reward level and required effort level at the time of the instruction cue and also the number of button presses actually made by the third person (Fig. 2B). A paired-samples t test revealed that subjects' overall task accuracy (mean = 93.93%) was significantly better than chance ($t_{(14)} = 54.5$, $p < 0.0001$). On the catch trials, the accuracy (mean = 78.64%) was also significantly greater than chance ($t_{(14)} = 12.76$; $p < 0.001$). These results indicate that subjects were attending to the reward value and the level of effort at the time of the instruction cues and also the number of button presses actually made by the third person. However, it was also important to demonstrate that all subjects' performances were similar for each level of effort and reward on the third-person trials to ensure that the instruction cues for the effort and reward levels were not acting as first-person risk or negative reward predictors on the third-person trials. To test this possibility, we performed a 2 × 2 repeated-measures ANOVA with the first factor being the reward level and the second being the amount of effort. There was no main effect of effort ($F_{(2.41, 31.31)} = 2.579$, $p = 0.083$), reward ($F_{(1,13)} = 0.024$, $p = 0.880$), and no interaction ($F_{(2.18, 28.32)} = 1.543$, $p = 0.231$) on task accuracy. This would suggest that the level of reward, and

crucially the number of cancellations being made by the confederate, did not affect the subjects ability to monitor the responses or indicate what level of reward was going to be received.

We also performed a repeated-measures ANOVA to determine whether there was a significant effect of net value on accuracy by breaking down the accuracy into six net value conditions. We found no significant effect of net value on task accuracy ($F_{(2.431, 31.599)} = 2.571$, $p > 0.05$). Therefore, our behavioral data suggest that there is no effect of net value, reward, or effort on the ability to perform the judgment task. Therefore, there is no difference in the risk of losing money between each of the different conditions on the third-person trials. Therefore, the effect of predicting a loss of money on the third-person trials would not vary across conditions and therefore would not vary with social net value. Risk thus cannot account for the activity that we identified in the ACC_g. Importantly, the high level of performance of subjects across all third-person conditions also ensures that any activity that we identified cannot be related to the subject predicting a reward in the first-person condition and predicting a punishment in the third-person condition. The performance of subjects on both tasks indicates that they were processing the reward value and the effort level at the time of the instructions cues on both first-person and third-person trials.

Imaging results
This study tested two hypotheses about the processing of cost–benefit analyses in the ACC. First, activity in the ACC_g covaries with the net value of rewards to be received by a third person, signaling a cost–benefit analysis for others' rewards. Second, activity in the ACC_s covaries with the net value of rewards on first-person trials, signaling one's own cost–benefit analysis. We examined activity time locked to the instruction cues and performed a parametric analysis to examine activity that covaried with the net value of rewards on both the first-person and third-person trials. We also performed a 2 × 2 × 2 factorial analysis to confirm the results of the parametric analysis. The first factor was the agency (first or third person), the second was reward level (high or low), and the third was effort, which was split into low (2 and 3 button presses) and high (8 and 12 button presses) conditions.

ACC_g and the net value of rewards on third-person trials
To test our first hypothesis, we investigated whether activity in the ACC_g was scaled with the net value of rewards on the third-person trials. To constrain our search to a hypothesized area, a mask of the ACC_g was used as a small volume correction for multiple comparisons. This mask ensured that any activated voxel at the group level would be within the ACC_g in 80% of the subjects. To ensure that the voxels identified in this analysis showed a significant effect on the third-person trials, but not a significant effect of net value on the first-person trials, we excluded any voxels in which activity covaried with the net value of rewards at the time of the first-person instruction cues. The voxels were excluded at a more liberal threshold ($p < 0.05$, uncorrected) to be conservative about the specificity of any response in the ACC_g to third-person net value. Activity in a cluster in the ACC_g was found to vary with the net value parameter (4, 22, 20, $Z = 2.8$, $p < 0.05$ small volume correction; Fig. 3A–C) in the midcingulate cortex (MCC), putatively area 24b′. It is important to note that the parametric effects identified in this analysis are from regressors that are orthogonal to the other parameter in the analysis (i.e., the β-coefficient for net value is for a regressor that is orthogonalized with respect to effort and vice versa for the effort β-coefficient). Therefore, the absence of a parametric effect

of effort (Fig. 3B) reflects only the fact that the unique variance of the effort parameter cannot account for activity in the ACC_g, not that effort does not influence activity in this area. By showing that the net value parameter significantly explains a unique portion of the variance, which is parameterized as effort divided by reward, we have shown an effect of both effort and reward on ACC_g activity.

An overlapping cluster with the same peak voxel also showed a significant interaction among agency, effort, and reward in the factorial analysis ($Z = 3.09$, $p < 0.05$ small volume correction). An additional, overlapping cluster that a contained the peak voxel from the parametric analysis also showed a main effect of agency ($Z = 5.57$, $p < 0.05$ small volume correction), highlighting this area as being differentially sensitive to first-person and third-person information. A cluster that overlapped with the peak voxel from the parametric analysis showed a significant effect of HR versus LR on the third-person trials ($Z = 3.66$, $p < 0.05$ small volume correction). No activity in any part of the ACC_g was found to covary with net value parameter on the first-person trials and no voxels showed a main effect of effort or reward ($p > 0.05$, uncorrected). A whole-brain analysis did not identify any voxels outside the ACC_g that covaried with the net value of rewards on the third-person trials, when using a whole-brain correction for multiple comparisons, or when correcting around the coordinates of a previous study that investigated first-person net value processing (Croxson et al., 2009). In addition, we found a simple effect between the HR, LE (2 or 3 cancellations) and the LR, HE conditions (8 or 12 cancellations) in a cluster in the ACC_g that overlapped with that showing a parametric effect ($Z = 3.22$, $p < 0.05$ small volume correction). That is, we found a significant difference between the trials with the two highest net value levels and the two lowest net value levels. We found no differences between the HR, HE and the LR, LE conditions in the ACC_g, even at a lowered threshold ($p > 0.01$, uncorrected). Therefore, we found no difference between the conditions with the same net value in the ACC_g on the third-person trials. Therefore, we identified activity in a portion of the ACC_g was scaled with the net value of rewards that were to be received specifically by a third person.

In addition to the hypothesized activity in the ACC_g, a second cluster, in a more posterior portion of area 24b′, showed a significant two-way interaction between agency and reward in the factorial analysis (4, 16, 26; $p < 0.05$ small volume corrected). This cluster did not overlap with that, which showed an effect of net value. This effect was driven by a differential response between the HR and LR conditions on the third-person trials ($p < 0.05$ small volume corrected). Therefore, distinct portions of the ACC_g process the net value of rewards and the undiscounted magnitude of a reward.

ACC_s and first-person effort processing
To test our second hypothesis, we investigated whether activity in the ACC_s was scaled with the net value of rewards on the first-person trials. To constrain the search to our hypothesized area, a mask of the ACC_s was used as a small volume correction for multiple comparisons. This mask (see Materials and Methods) ensured that any activated voxel at the group level would be within the ACC_s in 80% of the subjects. To ensure that any identified voxel in this analysis showed a significant effect exclusively on the first-person trials, we excluded any voxels that covaried with the net value of rewards at the time of the third-person instruction cues. The voxels were excluded at a more liberal threshold ($p < 0.05$, uncorrected) to be conservative about the
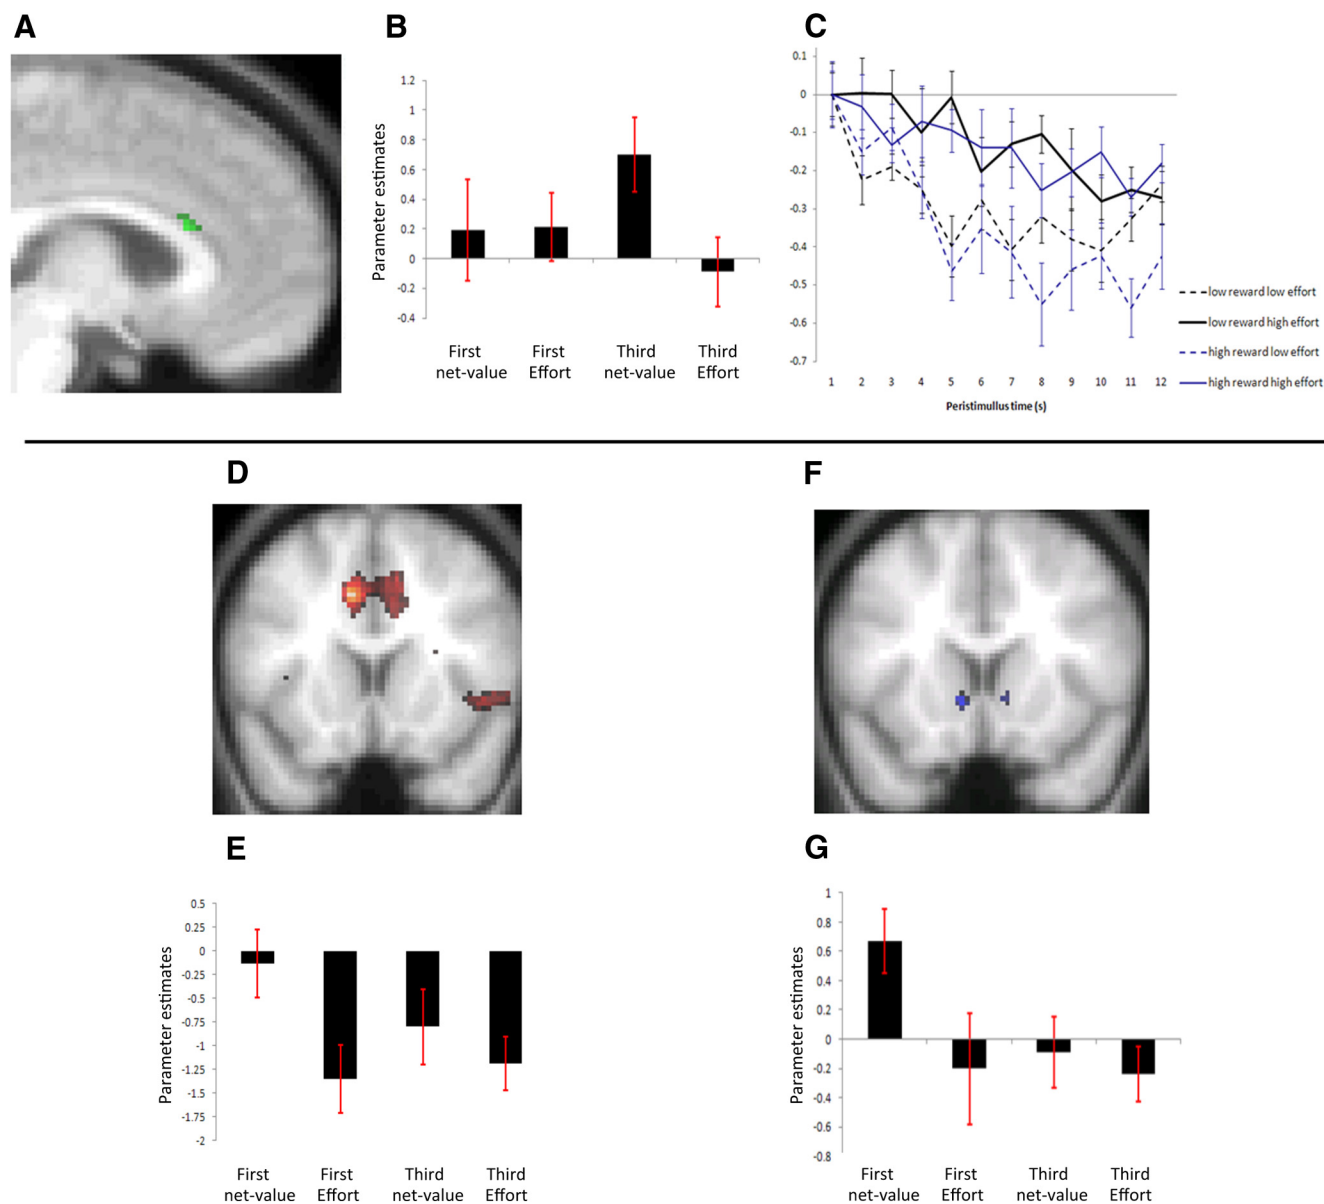
Figure 3. fMRI results. Activity shown in the ACC$_g$ (**A**) covaried with the net value of the rewards. **B**, **C**, β-coefficients (parameter estimates) (**B**) from the parametric analysis and perstimulus time histogram plots (**C**) of activity from the peak ACC$_g$ voxel. It is important to note that the β-coefficients from the parametric analysis reflect only the unique variance of a regressor. Therefore, the absence of a significant β-coefficient for the effort parameter in the ACC$_g$ suggests that activity in this region cannot be explained by the unique variance of the effort parameter. However, as can be seen in the PSTH plot, ACC$_g$ activity is modulated by both the effort level and the reward level. This supports the notion that activity in this portion of the ACC$_g$ covaries with the net value of another's reward. **D**, Activity shown in the ACC$_s$ that covaried with the effort level on both the first-person and third-person trials. **E**, Parameter estimates from the peak ACC$_s$ voxel. **F**, Activity shown in the NA that covaried with the net-value on the first-person trials only. **G**, Parameter estimates from the peak NA voxel. All error bars reflect SEM.

specificity of any response in the ACC$_g$ to first-person net value. There were no voxels in the ACC$_s$ in which activity covaried with net value on the first-person trials. There were also no voxels in the ACC$_s$ in which activity covaried with net value at the time of the third-person instruction cues.

We did not find activity in the ACC$_s$ to be scaled with the net value of rewards on first-person trials. However, previous studies have shown this area to be engaged during cost–benefit processing. Therefore, we performed further exploratory analyses to determine whether activity in the ACC$_s$ signaled any other information on the first-person trials. We found activity in a posterior portion of the ACC$_s$, in the MCC/rostral cingulate zone $(0, -22, 50; Z = 3.44, p < 0.05$ small volume corrected, putatively area 23c/24c$'$) that showed a main effect of effort (Fig. 3D). This

region therefore responded differentially to the level of effort regardless of the level of reward or agency. Notably, the profile of the ACC$_s$ response was consistent with some (Croxson et al., 2009), but not all previous studies examining activity at the time of cues that signal the level of reward and the amount of effort required. Our study found a negative relationship between effort and the BOLD response in the ACC$_s$, whereas other studies have shown a positive relationship (Prévost et al., 2010; Burke et al., 2013). This discrepancy may be explained by the fact that in this study and in Croxson et al. (2009), subjects were not engaged in deciding between differently valued options, as they were in the other studies. It is well known that activity in the ACC$_s$ is modulated during decision making by both chosen and unchosen options (Kolling et al., 2012). It is therefore possible that these

differences can be accounted for by subjects being able to make choices to minimize costs and maximize rewards during decision-making tasks, but only process the discounted value of the reward when effort and reward are instrumentally instructed, as in our study and that of Croxson et al. (2009). Therefore, our results still support previous accounts of this region's involvement in the processing of cost–benefit information, although we did not find that this region showed a sensitivity to reward magnitude

First-person net value
Although we did not find the hypothesized activity in the ACC$_s$, activity in the nucleus accumbens (NA) ($-8, 14, -4, Z = 2.8; p <$ 0.05 small volume corrected around the peak coordinate from Croxson et al., 2009) was found to covary with net value parameter at the time of the first-person instruction cues (Fig. 3). Our results therefore support previous findings that highlight the NA in cost–benefit-related information processing (Botvinick et al., 2009; Croxson et al., 2009; Ghods-Sharifi and Floresco, 2010; Day et al., 2011)

Discussion
We tested two hypotheses about the role of the ACC in processing the net value of rewards at the time that cues signaled the costs associated with rewards that would be received oneself or by a third person. Consistent with our first hypothesis, activity in the ACC$_g$ covaried with the net value of rewards to be received by the third person when the third person incurred the cost of the effort. Our second hypothesis, that activity in the ACC$_s$ would vary with the net value of rewards at the time of instruction cues on the first-person trials, was not supported. Rather, this region showed an effect of the anticipated level of effort on both the first-person and third-person trials. The ACC$_s$ signaled the effort level regardless of whether the effort was exerted oneself or by a third person. In addition, we found that activity in the NA scaled with the net value of rewards on first-person trials. Therefore, although the ACC$_s$ processed information about the costs associated with a reward, regardless of who worked to receive it, the ACC$_g$ was engaged when weighing up the benefits and costs associated with rewards that others were to receive.

Previous studies have suggested that the ACC$_s$ processes information about one's own decisions in a manner that conforms to the principles of reinforcement learning theory (Behrens et al., 2009), in which choices are made based on the predicted net value of decision-making outcomes. When an outcome is unexpected, prediction error signals code for the surprise evoked by the outcome, which serves to update future predictions (Rescorla and Wagner, 1972; Sutton and Barto, 1998). Neurophysiological and neuroimaging studies have identified activity in the ACC$_s$ that reflects the predicted net value of rewarding stimuli for oneself (Sallet et al., 2007; Quilodran et al., 2008; Jocham et al., 2009; Kennerley et al., 2009) and also neurons that signal that the value of an outcome is unexpectedly different from the predicted value (Amiez et al., 2005; Matsumoto et al., 2007; Sallet et al., 2007; Kennerley et al., 2011; Ribas-Fernandes et al., 2011). This evidence suggests that an important functional property of the ACC$_s$ is to signal predictions about the outcome of one's own decisions and to signal when they are discrepant from one's expectations (Alexander and Brown, 2011; Silvetti et al., 2014).

Neuroimaging studies suggest that the ACC$_g$ may mirror this property by showing that this area is activated when monitoring the unexpected outcomes of others' decisions (Apps et al., 2012; Apps et al., 2013a). One recent neurophysiology study found that

the ACC$_g$, and not the ACC$_s$, contains neurons that respond when a monkey is anticipating the delivery of a reward to another monkey (Chang et al., 2013). This suggests that the ACC$_g$ processes information about upcoming rewards that others will receive and that it is activated when the outcome of another's choice is unexpected. However, in these studies, there were no costs associated with the reward being delivered to another. Our study provides the first evidence that the human ACC$_g$ processes the predicted value of a reward that another will receive, supporting the claim that the ACC$_g$ processes information about others' rewards.

Anatomical evidence also supports the notion that the ACC$_g$ is engaged by both social and cost–benefit-related information. In monkeys, the homologous portion of the ACC$_g$ that was activated in this study (in the MCC) has strong connections to the posterior portions of the superior temporal sulcus, the temporal poles (Markowitsch et al., 1985; Seltzer and Pandya, 1989; Barbas et al., 1999), and the paracingulate cortex (Pandya et al., 1981; Vogt and Pandya, 1987; Petrides and Pandya, 2007). These three regions are believed to form a core circuit that is engaged when processing information about the mental states of others (Ramnani and Miall, 2004; Frith and Frith, 2006; Hampton et al., 2008). There is no evidence of connections between these regions and the ACC$_s$, supporting the notion that the information processed in the ACC$_g$ is more strongly linked to social behavior than that which is processed in the ACC$_s$. However, there is evidence to suggest that the ACC$_g$ is connected to the ACC$_s$ and the ventral striatum/NA. These two regions form a loop that is closed by return connections via the ventral pallidum and the thalamus (Groenewegen et al., 1993; Kunishio and Haber, 1994; Spooren et al., 1996; Middleton and Strick, 2000; Nakano et al., 2000; Haber and Knutson, 2010). It has been argued that this circuit is important for cost–benefit information processing. Disruptions of the striatopallidal connection (Mingote et al., 2008) and also to the striatocingulate connection (Hauber and Sommer, 2009) perturb normal behavioral patterns on tasks that require choices between options that have different associated costs. The connections of the ACC$_g$ with these regions suggests that the ACC$_g$ has access to information about the net value of rewarding outcomes. Therefore, anatomical evidence is consistent with the view that the ACC$_g$ processes net value when it relates to a reward another will receive.

Several lines of evidence support the notion that the ACC$_s$ and the NA are engaged during first person cost–benefit decision making. Lesions of the ACC$_s$ disrupt decision making on cost–benefit tasks (Walton et al., 2006; Hauber and Sommer, 2009). Single-unit recordings from neurons in this region in monkeys (Kennerley et al., 2009) and from homologous areas in rats (Hillman and Bilkey, 2010) have identified neurons in which spike frequency is a function of both the magnitude of a reward and the number of lever presses required for receipt. Furthermore, neurons in this region in rats signal the value of rewards discounted by the costs associated with social interaction (Hillman and Bilkey, 2012). Depletions of dopamine in the NA modulate cost–benefit-based decision making and neurophysiological recordings have identified neurons in the NA that show differential spike frequency related to high and LE conditions (Salamone et al., 2007; Font et al., 2008; Hauber and Sommer, 2009; Walton et al., 2009; Gan et al., 2010; Wanat et al., 2010). Similarly, previous fMRI studies have shown that activity in the ACC$_s$ and the NA is a function of the number of actions (Croxson et al., 2009; Kurniawan et al., 2010; Prévost et al., 2010) or the amount of cognitive effort (Botvinick et al., 2009; Schmidt et al., 2012) that has to

be exerted, suggesting that both regions may play an important role in signaling the effortful costs associated with choosing a rewarding option. The results of our study are broadly consistent with these findings, highlighting that activity in the NA signals the net value of first person rewards and activity in the ACC_s signals the effort-related costs associated with rewards that either oneself or another can obtain.

Our study is the first to examine activity at the time of cues that signaled the net value of rewards to be received by another. As a result, we were able to extend upon the findings of previous neuroimaging studies investigating the functional properties of the ACC_s in effort discounting and social decision making (Behrens et al., 2008; Prévost et al., 2010; Schmidt et al., 2012; Kurniawan et al., 2013; Meyniel et al., 2013). We show that this region processes the costs associated with both one's own and others' rewards, and not just the costs that will be incurred oneself. This finding is consistent with recent single-unit recording studies that identified ACC_s neurons that respond to both one's own or another's decision-making outcomes (Yoshida et al., 2012) and with a large corpus of studies that highlight the region as being engaged during empathic processing (i.e., during the processing one's own or another's pain) (Singer et al., 2005; Lockwood et al., 2013). However, it is notable that our results are not consistent with those of Croxson et al. (2009), who found a sensitivity to both effort and reward magnitude in the ACC_s. The absence of such an effect in our study could be due to the fact that that the effort period was a fixed time window in this study, whereas in Croxson et al. (2009), the duration of the effort period was dependent on the rate at which cancellations were made. It could therefore be argued that ACC_s activity may not be sensitive to the effort-discounted value of a reward, but to temporally discounted reward values. However, such an interpretation is inconsistent with a previous study showing that activity in the ACC_s is not sensitive to the temporally discounted reward values, but is sensitive to the effort-discounted value of rewards (Prévost et al., 2010). Therefore, the absence of an effect of net value in the ACC_s in this study is likely to be due to differences in the reward magnitudes between this and previous studies and not due to an insensitivity of ACC_s activity to rewards in general.

In summary, this study investigated the role of the ACC in processing cost–benefit analyses on rewards to be received oneself and by others. Our results highlight the ACC as an important structure in processing the net value of rewards that will be received by others and also in processing the costs associated with rewards regardless of who will receive them. However, these two functions are supported by the ACC_g and the ACC_s, respectively. This study further illuminates the important role that the ACC plays in processing information about others' decisions.

References

Alexander WH, Brown JW (2011) Medial prefrontal cortex as an action-outcome predictor. Nat Neurosci 14:1338–1344. CrossRef Medline

Amiez C, Joseph JP, Procyk E (2005) Anterior cingulate error-related activity is modulated by predicted reward. Eur J Neurosci 21:3447–3452. CrossRef Medline

Andersson JL, Hutton C, Ashburner J, Turner R, Friston K (2001) Modeling geometric deformations in EPI time series. Neuroimage 13:903–919. CrossRef Medline

Apps MA, Balsters JH, Ramnani N (2012) The anterior cingulate cortex: Monitoring the outcomes of others' decisions. Soc Neurosci 7:424–435. CrossRef Medline

Apps MA, Green R, Ramnani N (2013a) Reinforcement learning signals in the anterior cingulate cortex code for others' false beliefs. Neuroimage 64:1–9. CrossRef Medline

Apps MA, Lockwood PL, Balsters JH (2013b) The role of the midcingulate cortex in monitoring others' decisions. Front Neurosci 7:251. CrossRef Medline

Ashburner J, Friston KJ (2005) Unified segmentation. Neuroimage 26:839–851. CrossRef Medline

Barbas H, Ghashghaei H, Dombrowski SM, Rempel-Clower NL (1999) Medial prefrontal cortices are unified by common connections with superior temporal cortices and distinguished by input from memory-related areas in the rhesus monkey. J Comp Neurol 410:343–367. CrossRef Medline

Bautista LM, Tinbergen J, Wiersma P, Kacelnik A (1998) Optimal foraging and beyond: how starlings cope with changes in food availability. Am Nat 152:543–561. CrossRef Medline

Bautista LM, Tinbergen J, Kacelnik A (2001) To walk or to fly? How birds choose among foraging modes. Proc Natl Acad Sci U S A 98:1089–1094. CrossRef Medline

Behrens TE, Hunt LT, Woolrich MW, Rushworth MF (2008) Associative learning of social value. Nature 456:245–249. CrossRef Medline

Behrens TE, Hunt LT, Rushworth MF (2009) The computation of social behavior. Science 324:1160–1164. CrossRef Medline

Botvinick MM, Huffstetler S, McGuire JT (2009) Effort discounting in human nucleus accumbens. Cogn Affect Behav Neurosci 9:16–27. CrossRef Medline

Burke CJ, Brünger C, Kahnt T, Park SQ, Tobler PN (2013) Neural integration of risk and effort costs by the frontal pole: only upon request. J Neurosci 33:1706-+. CrossRef Medline

Chang SW, Gariépy JF, Platt ML (2013) Neuronal reference frames for social decisions in primate frontal cortex. Nat Neurosci 16:243–250. CrossRef Medline

Charnov EL (1976) Optimal foraging, marginal value theorem. Theor Popul Biol 9:129–136. CrossRef Medline

Croxson PL, Walton ME, O'Reilly JX, Behrens TE, Rushworth MF (2009) Effort-based cost-benefit valuation and the human brain. J Neurosci 29:4531–4541. CrossRef Medline

Day JJ, Jones JL, Wightman RM, Carelli RM (2010) Phasic nucleus accumbens dopamine release encodes effort- and delay-related costs. Biological Psychiatry 68:306–309. CrossRef Medline

Day JJ, Jones JL, Carelli RM (2011) Nucleus accumbens neurons encode predicted and ongoing reward costs in rats. Eur J Neurosci 33:308–321. CrossRef Medline

Deichmann R, Gottfried JA, Hutton C, Turner R (2003) Optimized EPI for fMRI studies of the orbitofrontal cortex. Neuroimage 19:430–441. CrossRef Medline

Floresco SB, Whelan JM (2009) Perturbations in different forms of cost/benefit decision making induced by repeated amphetamine exposure. Psychopharmacology 205:189–201. CrossRef Medline

Floresco SB, Tse MT, Ghods-Sharifi S (2008a) Dopaminergic and glutamatergic regulation of effort- and delay-based decision making. Neuropsychopharmacology 33:1966–1979. CrossRef Medline

Floresco SB, St Onge JR, Ghods-Sharifi S, Winstanley CA (2008b) Cortico-limbic-striatal circuits subserving different forms of cost-benefit decision making. Cogn Affect Behav Neurosci 8:375–389. CrossRef Medline

Font L, Mingote S, Farrar AM, Pereira M, Worden L, Stopper C, Port RG, Salamone JD (2008) Intra-accumbens injections of the adenosine A(2A) agonist CGS 21680 affect effort-related choice behavior in rats. Psychopharmacology 199:515–526. CrossRef Medline

Frith CD, Frith U (2006) The neural basis of mentalizing. Neuron 50:531–534. CrossRef Medline

Gan JO, Walton ME, Phillips PE (2010) Dissociable cost and benefit encoding of future rewards by mesolimbic dopamine. Nat Neurosci 13:25–27. CrossRef Medline

Ghods-Sharifi S, Floresco SB (2010) Differential effects on effort discounting induced by inactivations of the nucleus accumbens core or shell. Behav Neurosci 124:179–191. CrossRef Medline

Groenewegen HJ, Berendse HW, Haber SN (1993) organization of the output of the ventral striatopallidal system in the rat–ventral pallidal efferents. Neuroscience 57:113–142. CrossRef Medline

Haber SN, Knutson B (2010) The reward circuit: linking primate anatomy and human imaging. Neuropsychopharmacology 35:4–26. CrossRef Medline

Hampton AN, Bossaerts P, O'Doherty JP (2008) Neural correlates of mentalizing-related computations during strategic interactions in humans. Proc Natl Acad Sci U S A 105:6741–6746. CrossRef Medline

Hauber W, Sommer S (2009) Prefrontostriatal circuitry regulates effort-related decision making. Cereb Cortex 19:2240–2247. CrossRef Medline

Hillman KL, Bilkey DK (2010) Neurons in the rat anterior cingulate cortex dynamically encode cost-benefit in a spatial decision-making task. J Neurosci 30:7705–7713. CrossRef Medline

Hillman KL, Bilkey DK (2012) Neural encoding of competitive effort in the anterior cingulate cortex. Nat Neurosci 15:1290–1297. CrossRef Medline

Hutton C, Bork A, Josephs O, Deichmann R, Ashburner J, Turner R (2002) Image distortion correction in fMRI: a quantitative evaluation. Neuroimage 16:217–240. CrossRef Medline

Jocham G, Neumann J, Klein TA, Danielmeier C, Ullsperger M (2009) Adaptive coding of action values in the human rostral cingulate zone. J Neurosci 29:7489–7496. CrossRef Medline

Kagel JH, Levin D (1986) The winners curse and public information in common value auctions. American Economic Review 76:894–920.

Kennerley SW, Wallis JD (2009) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. Eur J Neurosci 29:2061–2073. CrossRef Medline

Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF (2006) Optimal decision making and the anterior cingulate cortex. Nat Neurosci 9:940–947. CrossRef Medline

Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. J Cogn Neurosci 21:1162–1178. CrossRef Medline

Kennerley SW, Behrens TE, Wallis JD (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. Nat Neurosci 14:1581-U1119. CrossRef Medline

Knutson B, Bossaerts P (2007) Neural antecedents of financial decisions. J Neurosci 27:8174–8177. CrossRef Medline

Kolling N, Behrens TE, Mars RB, Rushworth MF (2012) Neural mechanisms of foraging. Science 336:95–98. CrossRef Medline

Kunishio K, Haber SN (1994) Primate cingulostriatal projection- limbic striatal versus sensorimotor striatal input. J Comp Neurol 350:337–356. CrossRef Medline

Kurniawan IT, Seymour B, Talmi D, Yoshida W, Chater N, Dolan RJ (2010) Choosing to make an effort: the role of striatum in signaling physical effort of a chosen action. J Neurophysiol 104:313–321. CrossRef Medline

Kurniawan IT, Guitart-Masip M, Dayan P, Dolan RJ (2013) Effort and valuation in the brain: the effects of anticipation and execution. J Neurosci 33:6160–6169. CrossRef Medline

Lockwood PL, Sebastian CL, McCrory EJ, Hyde ZH, Gu X, De Brito SA, Viding E (2013) Association of callous traits with reduced neural response to others' pain in children with conduct problems. Curr Biol 23:901–905. CrossRef Medline

Markowitsch HJ, Emmans D, Irle E, Streicher M, Preilowski B (1985) Cortical and subcortical afferent connections of the primates temporal pole-a study of rhesus-monkeys, squirrel-monkeys, and marmosets. J Comp Neurol 242:425–458. CrossRef Medline

Matsumoto M, Matsumoto K, Abe H, Tanaka K (2007) Medial prefrontal cell activity signaling prediction errors of action values. Nat Neurosci 10:647–656. CrossRef Medline

Meyniel F, Sergent C, Rigoux L, Daunizeau J, Pessiglione M (2013) Neurocomputational account of how the human brain decides when to have a break. Proc Natl Acad Sci U S A 110:2641–2646. CrossRef Medline

Middleton FA, Strick PL (2000) Basal ganglia output and cognition: Evidence from anatomical, behavioral, and clinical studies. Brain Cogn 42:183–200. CrossRef Medline

Mingote S, Font L, Farrar AM, Vontell R, Worden LT, Stopper CM, Port RG, Sink KS, Bunce JG, Chrobak JJ, Salamone JD (2008) Nucleus accumbens adenosine A(2A) receptors regulate exertion of effort by acting on the ventral striatopallidal pathway. J Neurosci 28:9037–9046. CrossRef Medline

Nakano K, Kayahara T, Tsutsumi T, Ushiro H (2000) Neural circuits and functional organization of the striatum. J Neurol 247:V1–V15. CrossRef Medline

Nunes EJ, Randall PA, Santerre JL, Given AB, Sager TN, Correa M, Salamone JD (2010) Differential effects of selective adenosine antagonists on the effort-related impairments induced by dopamine d1 and d2 antagonism. Neuroscience 170:268–280. CrossRef Medline

Pandya DN, Van Hoesen GW, Mesulam MM (1981) efferent connections of the cingulate gyrus in the rhesus-monkey. Exp Brain Res 42:319–330. Medline

Petrides M, Pandya DN (2007) Efferent association pathways from the ros-

tral prefrontal cortex in the macaque monkey. J Neurosci 27:11573–11586. CrossRef Medline

Phillips PE, Walton ME, Jhou TC (2007) Calculating utility: preclinical evidence for cost-benefit analysis by mesolimbic dopamine. Psychopharmacology 191:483–495. CrossRef Medline

Prévost C, Pessiglione M, Météreau E, Cléry-Melin ML, Dreher JC (2010) Separate valuation subsystems for delay and effort decision costs. J Neurosci 30:14080–14090. CrossRef Medline

Quilodran R, Rothé M, Procyk E (2008) Behavioral shifts and action valuation in the anterior cingulate cortex. Neuron 57:314–325. CrossRef Medline

Ramnani N, Miall RC (2004) A system in the human brain for predicting the actions of others. Nat Neurosci 7:85–90. CrossRef Medline

Rescorla RA, Wagner AR (1972) Classical conditioning II: Current research and theory, pp 64–99. New York: Appleton-Century Crofts.

Ribas-Fernandes JJ, Solway A, Diuk C, McGuire JT, Barto AG, Niv Y, Botvinick MM (2011) A neural signature of hierarchical reinforcement learning. Neuron 71:370–379. CrossRef Medline

Rogers RD, Ramnani N, Mackay C, Wilson JL, Jezzard P, Carter CS, Smith SM (2004) Distinct portions of anterior cingulate cortex and medial prefrontal cortex are activated by reward processing in separable phases of decision-making cognition. Biol Psychiatry 55:594–602. CrossRef Medline

Rudebeck PH, Buckley MJ, Walton ME, Rushworth MF (2006) A role for the macaque anterior cingulate gyrus in social valuation. Science 313:1310–1312. CrossRef Medline

Salamone JD, Correa M, Farrar A, Mingote SM (2007) Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. Psychopharmacology 191:461–482. CrossRef Medline

Sallet J, Quilodran R, Rothé M, Vezoli J, Joseph JP, Procyk E (2007) Expectations, gains, and losses in the anterior cingulate cortex. Cogn Affect Behav Neurosci 7:327–336. CrossRef Medline

Schmidt L, Lebreton M, Clery-Melin M-L, Daunizeau J, Pessiglione M (2012) Neural mechanisms underlying motivation of mental versus physical effort. PLoS Biol 10.

Seltzer B, Pandya DN (1989) frontal-lobe connections of the superior temporal sulcus in the rhesus-monkey. J Comp Neurol 281:97–113. CrossRef Medline

Silvetti M, Alexander W, Verguts T, Brown J (2013) From conflict management to reward-based decision making: actors and critics in primate medial frontal cortex. Neurosci Biobehav Rev. Advance online publication. Retrieved Nov. 15, 2013. doi:10.1016/j.neubiorev.2013.11.003. CrossRef Medline

Singer T, Seymour B, O'Doherty J, Kaube H, Dolan R, Frith CD (2004) Empathy for pain involves the affective but not sensory components of pain. Science 303:1157–1162. CrossRef Medline

Spooren WP, Lynd-Balta E, Mitchell S, Haber SN (1996) Ventral pallidostriatal pathway in the monkey: Evidence for modulation of basal ganglia circuits. J Comp Neurol 370:295–312. CrossRef Medline

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. Cambridge, MA: MIT.

Vogt BA, Pandya DN (1987) Cingulate cortex of the rhesus monkey. 2. Cortical afferents. J Comp Neurol 262:271–289. CrossRef Medline

Vogt BA, Nimchinsky EA, Vogt LJ, Hof PR (1995) human cingulate cortex-surface-features, flat maps, and cytoarchitecture. J Comp Neurol 359:490–506. CrossRef Medline

Walton ME, Baudonnat M (2012) The value of competition in the rat race. Nat Neurosci 15:1182–1183. CrossRef Medline

Walton ME, Kennerley SW, Bannerman DM, Phillips PE, Rushworth MF (2006) Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. Neural Netw 19:1302–1314. CrossRef Medline

Walton ME, Groves J, Jennings KA, Croxson PL, Sharp T, Rushworth MF, Bannerman DM (2009) Comparing the role of the anterior cingulate cortex and 6-hydroxydopamine nucleus accumbens lesions on operant effort-based decision making. Eur J Neurosci 29:1678–1691. CrossRef Medline

Wanat MJ, Kuhnen CM, Phillips PE (2010) Delays conferred by escalating costs modulate dopamine release to rewards but not their predictors. J Neurosci 30:12020–12027. CrossRef Medline

Yoshida K, Saito N, Iriki A, Isoda M (2012) Social error monitoring in macaque frontal cortex. Nat Neurosci 15.